

APPARATUS AND METHOD TO PROVIDE INFORMATION TO MULTIPLE DATA STORAGE DEVICES

Field Of The Invention

This invention relates to an apparatus and method to provide information to
5 multiple data storage devices. More specifically, the invention is directed to storing
computer files on multiple information storage media.

Background Of The Invention

Many data processing systems require a large amount of data storage, for use in
efficiently accessing, modifying, and re-storing data. Data storage is typically separated
10 into several different levels, each level exhibiting a different data access time or data
storage cost. A first, or highest level of data storage involves electronic memory, usually
dynamic or static random access memory (DRAM or SRAM). Electronic memories take
the form of semiconductor integrated circuits where millions of bytes of data can be
stored on each circuit, with access to such bytes of data measured in nanoseconds. The
15 electronic memory provides the fastest access to data since access is entirely electronic.

A second level of data storage usually involves direct access storage devices
(DASD). DASD storage, for example, includes magnetic and/or optical disks. Data bits
are stored as micrometer-sized magnetically or optically altered spots on a disk surface,
representing the "ones" and "zeros" that comprise the binary value of the data bits.
20 Magnetic DASD includes one or more disks that are coated with remnant magnetic
material. The disks are rotatably mounted within a protected environment. Each disk is
divided into many concentric tracks, or closely spaced circles. The data is stored serially,

bit by bit, along each track. An access mechanism, known as a head disk assembly (HDA) typically includes one or more read/write heads, and is provided in each DASD for moving across the tracks to transfer the data to and from the surface of the disks as the disks are rotated past the read/write heads. DASDs can store gigabytes of data, and the

5 access to such data is typically measured in milliseconds (orders of magnitudes slower than electronic memory). Access to data stored on DASD is slower than electronic memory due to the need to physically position the disk and HDA to the desired data storage location.

Having a backup data copy is mandatory for many businesses for which data loss

10 would be catastrophic. The time required to recover lost data is also an important recovery consideration. With tape or library backup, primary data is periodically backed-up by making a copy on tape or library storage. One improvement over this arrangement is "flash copy," which mirrors contents of a primary device with a nearly identical secondary device. An example flash copy involves providing additional DASDs so that

15 data is written to the additional DASDs substantially in real time along with the primary DASDs. Then, if the primary DASDs fail, the secondary DASDs can be used to provide otherwise lost data. A drawback to this approach is that the number of required DASDs is doubled.

A different data backup alternative that avoids the need to provide double the

20 storage devices involves writing data to a redundant array of inexpensive devices (RAID). In this configuration, the data is apportioned among many DASDs. If a single

DASD fails, then the lost data can be recovered by applying error correction procedures to the remaining data. Several different RAID configurations are available.

The foregoing backup solutions are generally sufficient to recover data in the event that a storage device or medium fails. These backup methods are useful only for

- 5 device failures since the secondary data is a mirror of the primary data, that is, the secondary data has the same volume serial numbers (VOLSERs) and DASD addresses as the primary data. Data recovery due to system failures or storage controller failures, on the other hand, is not available using mirrored secondary data. Hence still further protection is required for recovering data if the entire system or even the site is destroyed
- 10 by a disaster such as an earthquake, fire, explosion, hurricane, etc. Disaster recovery requires that the secondary copy of data be stored at a location remote from the primary data. A known method of providing disaster protection is to periodically backup data to tape, such as a daily or weekly basis. The tape is then picked up by a vehicle and taken to a secure storage area usually located kilometers from the primary data location.
- 15 Nonetheless, this backup plan has its problems. First, it may take days to retrieve the backup data, and additional data is lost waiting for the backup data to be recovered. Furthermore, the same disaster may also destroy the storage location. A slightly improved backup method transmits data to a backup location each night. This allows the data to be stored at a more remote location. Again, some data may be lost between backups since
- 20 backups do not occur continuously, as in the dual copy solution. Hence, a substantial amount of data may still be lost and this may be unacceptable to some users.

More recently introduced data disaster recovery solutions include "remote dual copy," where data is backed-up not only remotely, but also continuously (either synchronously or asynchronously). In order to communicate duplexed data from one host processor to another host processor, or from one storage controller to another storage controller, or some combination thereof, a substantial amount of control data is required for realizing the process. A high overhead, however, can interfere with a secondary site's ability to keep up with a primary site's processing, thus threatening the ability of the secondary site to be able to recover the primary in the event a disaster occurs.

Disaster recovery protection for the typical data processing system requires that primary data stored on primary DASDs be backed-up at a secondary or remote location. The physical distance separating the primary and secondary locations can be set depending upon the level of risk acceptable to the user, and can vary from several kilometers to thousands of kilometers. The secondary or remote location, in addition to providing a backup data copy, must also have enough system information to take over processing for the primary system should the primary system become disabled. This is due in part because a single storage controller does not write data to both primary and secondary DASD storage at the primary and secondary sites. Instead, the primary data is stored on a primary DASD storage attached to a primary storage controller while the secondary data is stored on a secondary DASD storage attached to a secondary storage controller.

The secondary site must not only be sufficiently remote from the primary site, but must also be able to backup primary data in real time. The secondary site needs to

backup primary data in real time as the primary data is updated, with some minimal delay. Additionally, the secondary site has to backup the primary data regardless of the application program (e.g., IMS, DB2) running at the primary site and generating the data and/or updates. A difficult task required of the secondary site is that the secondary data

5 must be "order consistent," that is, secondary data is copied in the same sequential order as the primary data (sequential consistency) which requires substantial system considerations. Sequential consistency is complicated by the existence of multiple storage controllers each controlling multiple DASDs in a data processing system.

Without sequential consistency, secondary data inconsistent with primary data would

10 result, thus corrupting disaster recovery.

Remote data duplexing falls into two general categories, synchronous and asynchronous. Synchronous remote copy involves sending primary data to the secondary location and confirming the reception of such data before ending a primary DASD input/output (I/O) operation. Synchronous copy, therefore, slows the primary DASD I/O

15 response time while waiting for secondary confirmation. Primary I/O response delay is increased proportionately with the distance between the primary and secondary systems, a factor that limits the remote distance to tens of kilometers. Synchronous copy, however, provides sequentially consistent data at the secondary site with relatively little system overhead.

20 In contrast to synchronous remote copy, asynchronous remote copy provides better primary application system performance because the primary DASD I/O operation is completed without waiting for data to be confirmed at the secondary site. Therefore,

the primary DASD I/O response time is not dependent upon the distance to the secondary site and the secondary site can be thousands of kilometers remote from the primary site.

What is needed is a method that can implement a synchronous copy mode where that mode is required, but which can implement an asynchronous copy mode otherwise.

5 Applicants' invention comprises such a method.

Summary Of The Invention

Applicants' invention includes a method to provide information to multiple data storage devices. The method provides a first data storage device and a second data storage device, where the first data storage device is capable of communicating with the 10 second data storage device. The method provides first information and second information to the first data storage device at a first time. In certain embodiments, the second information comprises a write command. The method determines if the second information comprises a synchronous copy attribute.

If the second information comprises a synchronous copy attribute, then the 15 method synchronously provides the first information to the second data storage device. Alternatively, if the second information does not comprise a synchronous copy attribute, then the method provides the first information to the second data storage device at a second time, where that second time is later than the first time.

Brief Description Of The Drawings

20 The invention will be better understood from a reading of the following detailed description taken in conjunction with the drawings in which like reference designators are used to designate like elements, and in which:

FIG. 1 is a block diagram showing the components of Applicants' data storage and retrieval system;

FIG. 2 is a block diagram showing a device adapter in combination with two RAID ranks;

5 FIG. 3 is a block diagram showing a host computer, a first data storage device, and a second data storage device;

FIG. 4A is a flow chart summarizing certain steps in Applicants' method; and

FIG. 4B is a flow chart summarizing additional steps in Applicants' method.

Detailed Description Of The Preferred Embodiments

10 This invention is described in preferred embodiments in the following description with reference to the Figures, in which like numbers represent the same or similar elements. Referring now to FIG. 1. Applicants' information storage and retrieval system 100 includes a first cluster 101A and a second cluster 101B. Each cluster includes a processor portion 130 / 140 and an input/output portion 160 / 170. Internal PCI buses in 15 each cluster are connected via a Remote I/O bridge 155 / 165 between the processor portions 130 / 140 and I/O portions 160 / 170, respectively.

Information storage and retrieval system 100 further includes a plurality of host adapters 102 - 105, 107 - 110, 112 - 115, and 117 - 120, disposed in four host bays 101, 106, 111, and 116. Each host adapter may comprise one Fibre Channel port, one FICON 20 port, two ESCON ports, or two SCSI ports. Each host adapter is connected to both clusters through one or more Common Platform Interconnect buses 121 and 150 such that each cluster can handle I/O from any host adapter.

Processor portion 130 includes processor 132 and cache 134. In certain embodiments, processor 132 comprises a 64-bit RISC based symmetric multiprocessor. In certain embodiments, processor 132 includes built-in fault and error-correction functions. Cache 134 is used to store both read and write data to improve performance to 5 the attached host systems. In certain embodiments, cache 134 comprises about 4 gigabytes. In certain embodiments, cache 134 comprises about 8 gigabytes. In certain embodiments, cache 134 comprises about 12 gigabytes. In certain embodiments, cache 144 comprises about 16 gigabytes. In certain embodiments, cache 134 comprises about 32 gigabytes.

10 Processor portion 140 includes processor 142 and cache 144. In certain embodiments, processor 142 comprises a 64-bit RISC based symmetric multiprocessor. In certain embodiments, processor 142 includes built-in fault and error-correction functions. Cache 144 is used to store both read and write data to improve performance to the attached host systems. In certain embodiments, cache 144 comprises about 4 15 gigabytes. In certain embodiments, cache 144 comprises about 8 gigabytes. In certain embodiments, cache 144 comprises about 12 gigabytes. In certain embodiments, cache 144 comprises about 16 gigabytes. In certain embodiments, cache 144 comprises about 32 gigabytes.

I/O portion 160 includes non-volatile storage (“NVS”) 162 and NVS batteries 20 164. NVS 162 is used to store a second copy of write data to ensure data integrity should there be a power failure of a cluster failure and the cache copy of that data is lost. NVS 162 stores write data provided to cluster 101B. In certain embodiments, NVS 162

comprises about 1 gigabyte of storage. In certain embodiments, NVS 162 comprises four separate memory cards. In certain embodiments, each pair of NVS cards has a battery-powered charging system that protects data even if power is lost on the entire system for up to 72 hours.

5 I/O portion 170 includes NVS 172 and NVS batteries 174. NVS 172 stores write data provided to cluster 101A. In certain embodiments, NVS 172 comprises about 1 gigabyte of storage. In certain embodiments, NVS 172 comprises four separate memory cards. In certain embodiments, each pair of NVS cards has a battery-powered charging system that protects data even if power is lost on the entire system for up to 72 hours.

10 In the event of a failure of cluster 101B, the write data for the failed cluster will reside in the NVS 162 disposed in the surviving cluster 101A. This write data is then destaged at high priority to the RAID ranks. At the same time, the surviving cluster 101A will begin using NVS 162 for its own write data thereby ensuring that two copies of write data are still maintained.

15 I/O portion 160 further comprises a plurality of device adapters, such as device adapters 165, 166, 167, and 168, and sixteen disk drives organized into two RAID ranks, namely RAID rank “A” and RAID rank “B”. In certain embodiments, RAID ranks “A” and “B” utilize a RAID 5 protocol. In certain embodiments, RAID ranks “A” and “B” utilize a RAID 10 protocol.

20 Referring now to FIG. 2, SSA is a full-duplex loop-based architecture, with two physical read paths, i.e. read paths 240 and 250, and two physical write paths, i.e. write paths 260 and 270, to every disk drive attached to the loop. Device adapter 210 is

interconnected with RAID rank 220 and RAID rank 230 by loop 290. Data is sent from adapter card 210 to the first disk drive on the loop and then passed around the loop by the disk drives until it arrives at a target disk. Unlike bus-based architectures which reserve the entire bus for data transfer, SSA uses only the part of the loop between adjacent disk 5 drives for data transfer. Therefore, many simultaneous data transfers can take place on an SSA loop.

Applicants' invention includes a method to copy information to two or more information storage media. Applicants' method comprises what is sometimes called a "peer to peer remote copy," or PPRC, operation. Referring now to FIGs. 3 and 4A, in 10 step 510 Applicants' method generates one or more computer files, i.e. collectively "first information."

In certain embodiments, Applicants' first information comprises all or portions of one or more database files. In certain embodiments, Applicants' first information comprises one or more updates for one or more database files. In certain embodiments, 15 Applicants' first information comprises one or more log files, wherein those one or more log file comprise information relating to one or more database files. In certain embodiments, Applicants' first information comprises updates for one or more log files.

In certain embodiments, step 510 is performed by a computer, such as computer 390. In certain embodiments, computer 390 comprises a personal computer. In certain 20 embodiments, computer 390 comprises a work station. In certain embodiments, computer 390 comprises a main frame computer. In certain embodiments, computer 390 comprises a computer system comprising two or more computers, wherein those two or

more computers include zero or more personal computers, zero or more work stations, and zero or more main frame computers. In certain embodiments, Applicants' method transitions from step 510 to step 520.

In certain embodiments, Applicants' method includes step 512 wherein

- 5 Applicants' method writes the first information to a memory device, such as memory device 394 disposed in computer 390. In certain embodiments, computer 390 includes a controller 392. In certain embodiments, step 512 is performed by controller 392. In certain embodiments, one or more data base files, one or more database updates, one or more log files, one or more log file updates, and combinations thereof are stored on
- 10 device 396.

In certain embodiments, Applicants' method includes step 514 wherein

Applicants' method determines if the first information should be immediately written to one or more information storage media. In certain embodiments, step 514 is performed by a controller, such as controller 392. In certain embodiments, step 514 includes using

- 15 least recently used ("LRU") file management protocol wherein one or more of the least recently used files are written to one or more data storage devices, such as data storage devices 301 and 302, and then removed from resident memory, such as memory device 394 and/or 396.

If Applicants' method determines in step 514 that the first information should be

- 20 immediately written to one or more information storage media, then Applicants' method transitions from step 514 to step 520. Alternatively, if Applicants' method determines that the first information should not be immediately written to one or more information

storage media, then Applicants' method transitions from step 514 to step 510 and continues.

In step 520, Applicants' method provides the first information to a first data storage device. In certain embodiments, the first data storage device comprises a data storage and retrieval system, such as system 100 (FIG. 1). In certain embodiments, the first data storage device comprises a first cluster, such as cluster 101A, disposed in a data storage and retrieval system, such as system 100. In certain embodiments, the first data storage device comprises a first RAID rank, such as RAID rank 180 (FIG. 1).

In the illustrated embodiment of FIG. 3, computer 390 provides the first information to data storage device 301 via communication link 398. In certain embodiments, data storage device 301 comprises all or part of an data storage and retrieval system, such as system 100.

In certain embodiments, communication link 398 is selected from a serial interconnection, such as an RS-232 cable or an RS-422 cable, an ethernet interconnection, a SCSI interconnection, a Fibre Channel interconnection, an ESCON interconnection, a FICON interconnection, a Local Area Network (LAN), a private Wide Area Network (WAN), a public wide area network, Storage Area Network (SAN), Transmission Control Protocol/Internet Protocol (TCP/IP), the Internet, and combinations thereof.

In certain embodiments, step 520 further includes providing the first information to an I/O adapter, such as host adapter 302. In certain embodiments, step 520 is

performed by computer 390. In certain embodiments, step 520 is performed by a controller disposed in computer 390, such as controller 392.

In certain embodiments, Applicants' method includes step 522 wherein Applicants' method determines if the first information must be synchronously copied to 5 two or more information storage media. In certain embodiments, step 522 is performed by computer 390. In certain embodiments, step 522 is performed by controller 392.

If Applicants' method determines in step 522 that the first information must be synchronously copied to two or more information storage media, then Applicants' method transitions from step 522 to step 524 wherein Applicants' method generates a 10 write command comprising a synchronous copy attribute. In certain embodiments, step 524 includes setting a synchronous copy bit. In certain embodiments, step 524 includes setting a synchronous copy flag. Applicants' method transitions from step 524 to step 530. In certain embodiments, step 524 is performed by computer 390. In certain embodiments, step 524 is performed by controller 392.

15 In certain embodiments, if Applicants' method determines in step 522 that the first information need not be synchronously copied to two or more information storage media, then Applicants' method transitions from step 522 to step 526 wherein Applicants' method generates a write command comprising an asynchronous copy attribute. In certain embodiments, step 526 includes setting an asynchronous copy bit. In 20 certain embodiments, step 526 includes setting an asynchronous copy flag. In certain embodiments, Applicants' method generates a write command in step 526 that includes no copy attribute. In certain embodiments, step 526 is performed by computer 390. In

certain embodiments, step 526 is performed by controller 392. Applicants' method transitions from step 526 to step 530.

In step 530, Applicants' method provides a write command to the first data storage device. In certain embodiments, step 530 is performed by computer 390. In 5 certain embodiments, step 530 is performed by a controller disposed in computer 390, such as controller 392. In certain embodiments, step 520 is performed prior to step 530. In certain embodiments, step 530 is performed prior to step 520. In certain embodiments, steps 520 and 530 are performed substantially synchronously.

In step 540, Applicants' method, at a first write time, writes the first information 10 to a first information storage medium disposed in the first data storage device. In certain embodiments, that first information storage medium comprises a data cache, such as cache 332. In certain embodiments, that first information storage medium comprises non-volatile storage ("NVS"), such as NVS 362. In certain embodiments, step 540 is performed by processor 332.

15 In step 550, Applicants' method determines if the write command comprises a synchronous copy attribute. In certain embodiments, step 540 is performed prior to step 550. In certain embodiments, steps 540 and 550 are performed substantially synchronously. In certain embodiments, step 550 is performed by the first data storage device, such as data storage device 301. In the certain embodiments, step 550 is performed by a PPRC manager, such as PPRC manager 340. In certain embodiments, 20 step 550 is performed by a PPRC controller, such as PPRC controller 342.

If Applicants' method determines in step 550 that the write command of step 530 comprises a synchronous copy attribute, then Applicants' method transitions from step 550 to step 555 wherein Applicants' method synchronously provides the first information to a second data storage device, such as data storage device 401. In certain embodiments, 5 the second data storage device comprises a second data storage and retrieval system, such as system 100. In certain embodiments, the second data storage device comprises a second cluster, such as cluster 101B, disposed in a first data storage and retrieval system.

In certain embodiments, the second data storage device is selected by the host computer, such as computer 390. In certain embodiments, the second data storage device 10 is selected by a controller, such as controller 392, disposed in computer 390. In certain embodiments, the second data storage device is selected by the first PPRC manager, such as PPRC manager 340. In certain embodiments, the second data storage device is selected by a controller, such as controller 342, disposed in PPRC manager 340. In certain embodiments, the second data storage device is determined by programmable logic disposed in PPRC manager 340, such as, for example, computer useable medium 344 having computer readable program code disposed therein. In certain embodiments, the second data storage device is determined by A computer code product, such as computer code product 346, disposed in PPRC manager 340.

In the illustrated embodiment, the first information is provided to second data 20 storage device 401 using communication link 399. In certain embodiments, communication link 399 is selected from a serial interconnection, such as an RS-232 cable or an RS-422 cable, an ethernet interconnection, a SCSI interconnection, a Fibre

Channel interconnection, an ESCON interconnection, a FICON interconnection, a Local Area Network (LAN), a private Wide Area Network (WAN), a public wide area network, Storage Area Network (SAN), Transmission Control Protocol/Internet Protocol (TCP/IP), the Internet, and combinations thereof.

5 By “synchronously providing,” Applicants’ mean immediately providing the first information to the second data storage device. In certain embodiments, step 555 includes providing the first information to the second data storage device while that first information is being written to the first information storage medium.

In certain embodiments, step 555 is performed by the first PPRC manager, such 10 as PPRC manager 340. In certain embodiments, step 555 is performed by a controller, such as controller 342, disposed in PPRC manager 340. In certain embodiments, step 555 further includes providing the first information to a second host adapter, such as adapter 310, and a third host adapter, such as adapter 402.

Applicants’ method transitions from step 555 to step 565 wherein Applicants’ 15 method begins writing the first information to a second information storage medium disposed in the second data storage device. In certain embodiments, that second information storage medium comprises a data cache, such as cache 432. In certain embodiments, that second information storage medium comprises non-volatile storage (“NVS”), such as NVS 462. In certain embodiments, step 565 is performed by processor 20 432. In certain embodiments, steps 540 and 565 are performed at the first write time.

Referring now to FIG. 4B, in step 575 Applicants’ method determines if the first information has been written to both the first information storage medium and to the

second information storage medium. In certain embodiments, step 575 is performed by processor 332. In certain embodiments, step 575 is performed by processor 432. In certain embodiments, step 575 is performed by PPRC controller 342. In certain embodiments, step 575 is performed by PPRC controller 442. In certain embodiments, 5 step 575 is performed by a combination of one or more of PPRC controller 342, PPRC controller 442, processor 332, and/or processor 432.

If Applicants' method determines in step 575 that the first information has not been written to both the first information storage medium and to the second information storage medium, then Applicants' method continues to monitor the status of the writing 10 of the first information to the first information storage medium and to the second information storage medium. If Applicants' method determines in step 575 that the first information has been written to both the first information storage medium and to the second information storage medium, then Applicants' method transitions from step 575 to step 585 wherein Applicants' method generates a signal that the write command is 15 complete.

In certain embodiments, step 585 further includes providing that write complete signal to computer 390. In these embodiments, the write complete signal of step 585 is provided to computer 390 at a synchronous write complete time.

In certain embodiments, step 585 is performed by processor 332. In certain 20 embodiments, step 585 is performed by processor 432. In certain embodiments, step 585 is performed by PPRC manager 340. In certain embodiments, step 585 is performed by PPRC controller 342. In certain embodiments, step 585 is performed by PPRC manager

440. In certain embodiments, step 585 is performed by PPRC controller 442. In certain embodiments, step 585 is performed by a combination of one or more of PPRC manager 340, PPRC controller 342, PPRC manager 440, PPRC controller 442, processor 332, and/or processor 432.

5 Referring now to FIGs. 4A and 4B, if Applicants' method determines in step 550 that the write command of step 530 does not comprise a synchronous copy attribute, then Applicants' method transitions from step 550 to step 560 wherein Applicants' method determines if the first information has been written to the first information storage medium. If Applicants' method determines in step 560 that the first information has not
10 been written to the first information storage medium, then Applicants' method continues to monitor the writing of the first information to the first information storage medium.

If Applicants' method determines in step 560 that the first information has been written to the first information storage medium, then Applicants' method transitions from step 560 to step 570 wherein Applicants' method generates a signal that the write
15 command is complete. In certain embodiments, step 570 is performed by processor 332. In certain embodiments, step 570 is performed by PPRC manager 340. In certain embodiments, step 570 is performed by PPRC controller 342.

In certain embodiments, step 570 further includes providing the write command complete signal to computer 390. In these embodiments, the write complete signal of
20 step 570 is provided to computer 390 at an asynchronous write complete time. In certain embodiments, the write command complete signal is provided to computer 390 by a host adapter, such as adapter 302.

In step 580, Applicants' method schedules the first information for later transmission to a second data storage device. In certain embodiments, step 580 is performed by processor 332. In certain embodiments, step 580 is performed by PPRC manager 340. In certain embodiments, step 580 is performed by PPRC controller 342. In 5 certain embodiments, the time to provide the first information to the second data storage device is disposed in programmable logic, such as logic 344. In certain embodiments, the time to provide the first information to the second data storage device is disposed in firmware, such as firmware 346.

In step 590, Applicants' method provides the first information to a second data 10 storage device. In certain embodiments, the second data storage device comprises a second data storage and retrieval system, such as system 100. In certain embodiments, the second data storage device comprises a second cluster in the same data storage and retrieval system, such as cluster 101B. In the illustrated embodiment of FIG. 3, the second data storage device comprises assembly 401.

15 In certain embodiments, the second data storage device is selected in step 590 by the host computer, such as computer 390. In certain embodiments, the second data storage device is selected in step 590 by a controller, such as controller 392, disposed in computer 390. In certain embodiments, the second data storage device is selected in step 590 by the first PPRC manager, such as PPRC manager 340. In certain embodiments, the 20 second data storage device is selected in step 590 by a controller, such as controller 342, disposed in PPRC manager 340. In certain embodiments, the second data storage device is determined by programmable logic, such as logic 344, disposed in PPRC manager 340.

In certain embodiments, the second data storage device is determined by firmware, such as firmware 346, disposed in PPRC manager 340.

In the illustrated embodiment, step 590 includes providing the first information to second data storage device 401 using communication link 399. In certain embodiments, 5 communication link 399 is selected from a serial interconnection, such as an RS-232 cable or an RS-422 cable, an ethernet interconnection, a SCSI interconnection, a Fibre Channel interconnection, an ESCON interconnection, a FICON interconnection, a Local Area Network (LAN), a private Wide Area Network (WAN), a public wide area network, Storage Area Network (SAN), Transmission Control Protocol/Internet Protocol (TCP/IP), 10 the Internet, and combinations thereof.

In certain embodiments, step 590 is performed by a data storage and retrieval system processor, such as processor 332. In certain embodiments, step 590 is performed by a PPRC manager, such as PPRC manager 340. In certain embodiments, step 590 is performed by a PPRC controller, such as controller 342.

15 In step 595, Applicants' method writes the first information to a second information storage medium at a second write time, where that second write time is later than the first write time of step 540 and where that second write time is later than the asynchronous write complete time of step 570. In certain embodiments, that second information storage medium comprises a data cache, such as cache 432. In certain 20 embodiments, that second information storage medium comprises non-volatile storage ("NVS"), such as NVS 462. In certain embodiments, step 565 is performed by processor

432. In certain embodiments, step 595 is performed by a data storage device processor, such as processor 432.

In certain embodiments, individual steps recited in FIGs. 4A and/or 4B may be combined, eliminated, or reordered.

5 Applicants' invention further includes an article of manufacture comprising a computer useable medium, such as for example computer useable media 344 and/or 444, having computer readable program code disposed therein method to copy information to two or more information storage media. In certain embodiments, the article of manufacture comprises a data storage and retrieval system. In certain embodiments, the
10 article of manufacture comprises a PPRC copy manager. In certain embodiments, the computer readable program code implements the steps of FIGs. 4A and 4B.

Applicants' invention further includes a computer program product, such as for example computer program product 346 and/or 446, usable with a programmable computer processor having computer readable program code embodied therein method to
15 provide information to two or more information storage media. In certain embodiments, the computer program product is disposed in a data storage and retrieval system. In certain embodiments, the computer program product is disposed in a processor disposed in a data storage and retrieval system. In certain embodiments, the computer program product is disposed in a PPRC manager. In certain embodiments, the computer program
20 product is disposed in a PPRC manager. In certain embodiments, the computer program code implements the steps of FIGs. 4A and 4B.

While the preferred embodiments of the present invention have been illustrated in detail, it should be apparent that modifications and adaptations to those embodiments

may occur to one skilled in the art without departing from the scope of the present invention as set forth in the following claims.